RESEARCH ARTICLE · COMPUTATIONAL AND SYSTEMS BIOLOGY · ONCOLOGY

Pairwise Co-expression Instability in Lung Adenocarcinoma: A Systematic Large-Scale Analysis of Tumour–Normal Differential Gene Pair Correlations and Identification of Candidate Two-Gene Network Vulnerabilities

Author: Janos Gabor Melegh

Submitted: March 2026

ABSTRACT

Background. Transcriptomic co-expression networks in lung adenocarcinoma (LUAD) are structurally distinct from those of normal lung tissue, exhibiting earlier loss of global connectivity and reduced algebraic cohesion. Single-gene perturbation approaches have not yielded consistently exploitable vulnerabilities, motivating a systematic shift towards pairwise relational analysis.

Methods. We performed a large-scale, batch-processed pairwise Pearson correlation analysis across a joint gene universe of 4,000 genes in 513 LUAD tumour and 288 matched normal samples from TCGA, computing differential co-expression (Δr) for all possible gene pairs (n = 7,998,000). Candidate pairs were prioritised through a hierarchical multi-layer framework comprising resampling robustness assessment (200 iterations, 80% subsampling), local network context analysis (top-25 neighbourhood Jaccard overlap), deep and ultra-deep stress testing (up to 3,500 iterations per pair), 800-gene subnetwork spectral analysis, and label-shuffled permutation testing (1,200 shuffles per pair). Evidence-integrated priority scoring incorporated internal network metrics alongside first-pass external evidence annotation.

Results. All 7,998,000 gene pairs were evaluated with verified run integrity (zero mismatch; zero validation errors across 320 batches). A primary candidate panel of 10 hub–hub gene pairs with uniformly tumour-strengthened correlations (mean Δr = 0.865; sign consistency = 1.00 across 200 resampling replicates) was identified. Local network context analysis revealed that primary pairs share substantial tumour-specific neighbourhood overlap (mean Jaccard = 0.568) that is absent in normal tissue (0.000), a pattern inverted in the 10-pair control panel. Label-shuffled permutation testing confirmed that primary pair signals are extreme relative to a null label distribution (null-extremeness scores: 8.1–36.3). Two architecturally distinct primary subtypes were identified: pair-dominant pairs (n = 4; focal relational breakpoints; 800-gene module mean $|\Delta r| \approx 0.15$–0.16) and module-dominant pairs (n = 6; broad local subnetwork reorganisation; module mean $|\Delta r| \approx 0.32$–0.36). An evidence-integrated priority framework yielded a final decision structure of one primary experimen-

tal candidate, two reserve candidates, and five currently deprioritised pairs, with the lead candidate demonstrating threshold-robust score separation (Δscore = 0.269 relative to the second-ranked pair).

Conclusions. This study provides a systematic, multi-layer validated catalogue of tumour-specific gene pair co-expression instability in LUAD and identifies a ranked panel of candidate dual-gene network vulnerabilities suitable for experimental follow-up. The computational framework described is scalable and reproducible, and may be applicable to other tumour types.

Keywords: lung adenocarcinoma; co-expression network; pairwise correlation; network instability; gene pair vulnerability; TCGA; computational oncology

1. INTRODUCTION

Lung adenocarcinoma (LUAD) is the most common histological subtype of non-small-cell lung cancer and remains among the leading causes of cancer-related mortality worldwide.[1] Despite advances in targeted therapy and immunotherapy, a substantial proportion of patients do not harbour readily actionable single-gene alterations, and durable responses to available agents remain limited in the majority of cases.[2] This has renewed interest in network-level approaches to target identification, grounded in the observation that tumour transcriptomes are characterised not merely by differential gene expression but by systematic reorganisation of co-regulatory architecture.[3,4]

Prior work from our group and others has established that LUAD tumour gene co-expression networks lose global connectivity at lower correlation thresholds than matched normal lung tissue.[5] Bootstrap resampling across 2,000-gene subsets confirmed that the tumour-specific fragmentation threshold (r^* at which the giant connected component fraction falls below 0.9) is consistently lower in tumour than in normal tissue in 85% of replicate pairs ($P = 0.85$), with tumour r^* values exhibiting greater dispersion ($\sigma \approx 0.031$ versus 0.013 in normal), consistent with increased transcriptional heterogeneity.[5] In breast cancer (TCGA-BRCA), parallel analyses using algebraic connectivity (λ_2) as the network cohesion metric revealed that tumour networks collapse to near-zero connectivity at a correlation threshold of $r \approx 0.110$, compared with $r \approx 0.250$ in normal tissue, yielding a 2.27-fold stability ratio and a 56% reduction in cohesion.[6] These findings collectively suggest that tumour co-expression architecture is systematically less stable than that of normal tissue, and that this instability may be exploitable as a therapeutic target.

Single-gene perturbation strategies directed at transcriptional network hubs have provided limited traction, in part because the hub identity of any gene is context-dependent and because loss of a single highly connected node is frequently buffered by network redundancy.[7] By contrast, the concept of two-gene co-expression vulnerability—in which the simultaneous disruption of a tumour-stabilised gene pair produces a disproportionate network effect relative to either perturbation alone—represents an underexplored therapeutic logic in LUAD. Synthetic lethality screens and combinatorial perturbation studies in other tumour types have demonstrated that pairwise vulnerabilities can exhibit strong tumour-selective phenotypes.[8]

We reasoned that a systematic, exhaustive analysis of differential pairwise co-expression across the LUAD tumour–normal transcriptomic space, coupled with multi-layer robustness validation and network context characterisation, could provide a ranked catalogue of candidate two-gene network vulnerabilities suitable for experimental prioritisation. Here, we describe such an analysis, encompassing 7,998,000 gene pairs evaluated across a validated 4,000-gene joint universe, and report the identification of a primary candidate panel with

consistent tumour-specific co-expression strengthening, tumour-localised local network coherence, and resistance to resampling and label-permutation-based null testing.

2. METHODS

2.1 Data sources and cohort definition

Gene expression data were obtained from The Cancer Genome Atlas (TCGA) LUAD cohort (TCGA-LUAD) via the Genomic Data Commons (GDC) portal. Samples were stratified by tissue type into a tumour group (n = 513) and a normal adjacent tissue group (n = 288). All expression values were log2-transformed TPM counts derived from RNA-seq (IlluminaHiSeq 2000 platform). Sample provenance, batch assignment and QC metrics were retrieved from TCGA clinical and auxiliary data files. This study involves secondary analysis of publicly available de-identified data and did not require ethics committee approval.

2.2 Gene universe construction and input preparation

A joint gene universe of 4,000 genes was constructed by selecting genes expressed in both tumour and normal sample sets (≥0.1 TPM in ≥10% of samples in each group), followed by ranking by mean absolute deviation across the combined sample pool and retaining the top 4,000 by variability. This procedure ensures that network comparisons between the two tissue groups are computed on identical gene sets, a prerequisite for valid network-level inference as demonstrated by simulation analysis.[5] All downstream analyses were performed on audited, rebuilt expression matrices derived from the canonical input set; legacy processed files were explicitly excluded from the production pipeline.

2.3 Pairwise co-expression analysis

Pearson correlation coefficients were computed for all C(4000, 2) = 7,998,000 unique gene pairs separately in tumour and normal expression matrices. The primary metric of interest was the signed differential correlation:

$$\Delta r = r_{tumour} - r_{normal}$$

and the absolute differential correlation $|\Delta r|$ as a direction-agnostic measure of co-expression change. Computation was parallelised across 320 batches (approximately 25,000 pairs per batch) using a checkpointed batch architecture in which each batch produced an independent result file and an associated curve file. Batch completion and file integrity were verified post-hoc through a dedicated integrity check module confirming batch count, result–curve file correspondence, pair count consistency, and absence of validation errors (run_integrity_ok = true; mismatches = 0; validation errors = 0).

2.4 Candidate pair prioritisation framework

Initial candidate selection employed three parallel shortlist strategies applied to the top 2,000 pairs by $|\Delta r|$: (i) top 100 pairs by absolute differential correlation; (ii) top 150 pairs by a composite score combining $|\Delta r|$ and gene-level recurrence frequency in the top 2,000; and (iii) top 100 pairs in which both member genes rank among the top 50 genes by recurrence (hub–hub intersection). The union of these shortlists was scored on a validation index incorporating shortlist membership count, hub status and directional consistency,

yielding a master review table from which a primary candidate panel of 10 pairs and a structural control panel of 10 pairs were selected.

2.5 Resampling robustness assessment

To assess the stability of candidate pair rankings under sampling variation, a resampling procedure was applied to all 20 candidate pairs. In each of 200 iterations, 80% of tumour samples (n = 410 from 513) and 80% of normal samples (n = 230 from 288) were drawn without replacement, Pearson correlations were recomputed, and Δr was calculated. Stability metrics included: mean and standard deviation of resampled Δr; sign consistency (proportion of iterations in which the sign of Δr agreed with the full-cohort estimate); direction agreement with the full-cohort Δr (direction_match); and a binary strict robustness flag (robustness_flag_strict = 1 if sign consistency = 1.00 and n_successful_reps = 200/200).

2.6 Local network context analysis

For each candidate pair (gene A, gene B), the top 25 co-expression partners of each gene were identified independently in the tumour and normal correlation matrices. Local network context metrics included: (i) tumour-versus-normal Jaccard similarity of the top-25 partner sets for gene A and gene B individually; (ii) pairwise neighbourhood overlap in tumour (Jaccard between the top-25 partner sets of A and B in the tumour matrix); (iii) pairwise neighbourhood overlap in normal; (iv) a context similarity shift (tumour minus normal overlap); and (v) a hub-like pair context flag (set to 1 if tumour overlap exceeded a predefined threshold and normal overlap fell below it, indicating tumour-specific local co-localisation).

2.7 Deep, ultra-deep, and mega-network stress testing

The 10 primary candidate pairs were subjected to progressively more stringent computational stress testing. *Deep validation* employed 1,000 resampling iterations per pair (sample fraction 0.80) with neighbourhood characterisation extended to the top 100 partners and shared context to the top 50, yielding bridge scores and shared-context shift metrics. *Ultra-deep validation* employed five subsampling fractions (0.50, 0.60, 0.70, 0.80, 0.90), 400 iterations per fraction (total 2,000 subsample replicates), 1,500 bootstrap replicates, 300 rewired-partner null iterations, and top-250 neighbourhood characterisation. *Mega-network stress testing* extended local subnetwork size to 500 genes per anchor, employed six subsampling fractions (0.40–0.90) with 180 iterations per fraction, and computed spectral metrics including the five leading eigenvalues of the local subnetwork Laplacian.

2.8 Label-shuffled permutation testing

To assess whether primary pair signals are extreme relative to a null distribution generated by random tumour–normal label assignment, a label-shuffle permutation procedure was applied. For each of the 10 primary pairs, 1,200 label-shuffled null iterations were computed over 800-gene local subnetworks centred on each pair's anchor genes (top 400 co-expression partners per anchor). In each iteration, tumour and normal sample labels were randomly permuted, correlation matrices recomputed, and local subnetwork metrics re-evaluated. Empirical null-extremeness scores (null_extremeness_score) were derived from the position of the observed statistics in the null distribution. Spectral instability metrics (Frobenius norm of tumour–normal subnetwork Laplacian difference, top-five eigenvalue dispersion) were computed for both observed and null

conditions. Total runtime for this stage was approximately 35.75 minutes on the production compute environment.

2.9 Multi-layer priority integration and evidence scoring

Results from deep, ultra-deep, mega-network, and label-shuffle runs were integrated into a null-integrated priority table. For each pair, pair-dominance class (pair-dominant, module-dominant, or mixed) was assigned based on the ratio of pair-level Δr to 800-gene module mean $|\Delta r|$ in the null-run output (pair-to-module ratio). Pairs with ratio ≥ 4.0 were classified as pair-dominant; pairs with ratio < 3.5 as module-dominant; intermediate pairs as mixed. A composite priority score was computed weighting pair-dominant score (from pair-level null metrics), module-dominant score (from subnetwork-level metrics), and overall null-extremeness. An evidence-integrated refreshed priority score additionally incorporated first-pass external annotation for LUAD biological relevance, dependency support, synthetic lethality or vulnerability association, novelty status, actionability, and evidence confidence level, each scored 0.0–1.0 and combined with weighted internal score normalisation.

2.10 Experimental design framework (Batch A)

For the eight-pair focus panel derived from the null-integrated priority analysis, a structured experimental design was generated specifying four perturbation arms per pair: gene-A-only perturbation, gene-B-only perturbation, A+B combined perturbation, and unperturbed control reference. Track-specific primary hypotheses and success criteria were defined separately for pair-dominant pairs (A+B superiority over single perturbations on focal pair-edge rewiring and tumour-selective functional impact) and module-dominant pairs (A+B superiority on local subnetwork-level rewiring metrics). Execution priority was assigned by tier (Tier 1, 2, or 3) based on the composite priority schedule.

3. RESULTS

3.1 Large-scale pairwise analysis: coverage and integrity

All 7,998,000 possible gene pairs within the 4,000-gene LUAD joint universe were successfully evaluated. The production pipeline completed 320/320 planned batches with verified integrity: result–curve file mismatch count = 0; validation error count = 0; pair count error count = 0; run_integrity_ok = true. The top 2,000 pairs by $|\Delta r|$ comprised 718 unique genes, indicating that the most extreme differential co-expression signals are concentrated within a non-trivial but circumscribed gene set, consistent with the existence of discrete network reorganisation foci rather than diffuse transcriptome-wide noise.

7,998,000	**320 / 320**	**718**	**0**
Gene pairs evaluated	Batches completed (integrity verified)	Unique genes in top 2,000 pairs	Integrity errors or mismatches

3.2 Primary and control candidate panels

Application of the hierarchical prioritisation framework yielded a primary candidate panel of 10 gene pairs and a structural control panel of 10 pairs (Table 1). All 10 primary pairs were classified as hub–hub (both

member genes ranking in the top recurrence stratum) and exhibited tumour-strengthened correlation direction ($\Delta r > 0$; tumour_stronger). The control panel comprised 9 nonhub pairs and one hub–partner pair, predominantly exhibiting normal-strengthened correlation direction (normal_stronger). This structural and directional dichotomy between panels was preserved throughout all downstream analyses.

Table 1. Summary characteristics of the primary and control candidate panels.

PANEL	N PAIRS	DOMINANT STRUCTURAL CLASS	DOMINANT DIRECTION	MEAN |ΔR| (FULL COHORT)	MEAN RESAMPLED ΔR (200 REP)	SIGN CONSISTENCY
Primary	10	hub–hub (10/10)	tumour_stronger (10/10)	0.872	0.865	1.000
Control	10	nonhub_pair (9/10)	normal_stronger (10/10)	0.793	−0.789	1.000

Sign consistency: proportion of resampling iterations in which the sign of Δr agreed with the full-cohort estimate. Resampling: 200 iterations, 80% subsampling fraction, Pearson correlation.

3.3 Resampling stability of candidate pairs

Resampling analysis (200 iterations; 80% subsampling fraction) confirmed full sign and direction stability for all 20 candidate pairs: n_successful_reps = 200/200, sign_consistency = 1.00, and direction_match_vs_full = 1.00 for every pair, yielding robustness_flag_strict = 1 universally. Panel-level statistics are summarised in Table 1. The primary panel showed a mean resampled $\Delta r = 0.865$ (SD = 0.034), with mean tumour-side correlation 0.852 and mean normal-side correlation −0.013, indicating that tumour co-expression of primary pairs approaches near-complete positive correlation while normal co-expression is negligible to weakly negative. Conversely, the control panel showed mean resampled $\Delta r = -0.789$ (SD = 0.030), with mean normal-side correlation 0.872 and tumour-side correlation 0.083, confirming the structurally inverted nature of control pair behaviour.

These findings indicate that the primary–control distinction is not an artefact of a single full-cohort measurement but is reproducible under repeated random subsampling, excluding trivial explanations based on sampling sensitivity or outlier-driven correlation inflation.

3.4 Tumour-specific local network context

Local network context analysis revealed a pronounced structural contrast between primary and control panels in their neighbourhood organisation within tumour versus normal co-expression matrices (Table 2). Primary pairs exhibited high tumour-specific pairwise neighbourhood overlap (mean Jaccard = 0.568) that was entirely absent in normal tissue (mean = 0.000), indicating that the co-expression partners of both member genes overlap substantially in the tumour environment but not in normal tissue. All 10 primary pairs met the hub-like pair context criterion (flag = 1), and the mean context similarity shift was +0.979. The control panel showed the inverse pattern: normal neighbourhood overlap (0.867) with near-zero tumour overlap (0.002), and a mean context shift of −0.821.

Table 2. Panel-level local network context metrics.

PANEL	TUMOUR PAIR NEIGHBOUR OVERLAP (MEAN JACCARD)	NORMAL PAIR NEIGHBOUR OVERLAP (MEAN JACCARD)	CONTEXT SIMILARITY SHIFT (TUMOUR − NORMAL)	HUB-LIKE PAIR CONTEXT FLAG (SUM)
Primary	0.568	0.000	+0.979	10 / 10
Control	0.002	0.867	−0.821	0 / 10

Neighbourhood overlap: Jaccard similarity of the top-25 co-expression partners of gene A and gene B within each matrix. Context similarity shift: tumour overlap minus normal overlap per pair, panel mean reported.

These results suggest that primary pairs are not merely isolated pairs of genes with incidentally altered pairwise correlation, but are embedded within tumour-specific local co-expression modules—in which both member genes acquire a shared neighbourhood—that dissolve in normal tissue. This property distinguishes these pairs from candidates selected on Δr magnitude alone and provides independent network-structural support for their biological specificity.

3.5 Two architecturally distinct primary pair subtypes

Mega-network stress testing (500-gene local subnetworks, spectral metrics) and label-shuffled null analysis (800-gene subnetworks, 1,200 permutations per pair) together enabled stratification of the 10 primary pairs into two architecturally distinct subtypes based on the ratio of pair-level Δr to 800-gene subnetwork mean $|\Delta r|$ (pair-to-module ratio; Table 3).

Table 3. Primary pair architectural subtype classification and key metrics.

SUBTYPE	N	PAIR-TO-MODULE RATIO (NULL RUN)	800-GENE MODULE MEAN |ΔR|	NULL-EXTREMENESS SCORE RANGE	NULL-STRENGTH (PAIR/MODULE)	INTERPRETATION
Pair-dominant	4	5.26–6.35	0.151–0.157	32.3–36.3	300 / 300 (ceiling)	Focal relational breakpoint; pair-level signal predominates over subnetwork-level reorganisation
Module-dominant	3	2.56–2.58	0.316–0.332	9.3–9.8	300 / 300 (ceiling)	Broad local subnetwork reorganisation; pair represents a nexus within a wider tumour-rewired module
Mixed	3	2.64–2.77	0.316–0.326	7.99–8.41	300 / 300 (ceiling)	Intermediate; both focal and subnetwork-level signals are present

Null-strength metrics (pair-level and module-level) reached the computational ceiling (300) for all 10 primary pairs, indicating that observed signals were never exceeded by the null distribution in 1,200 label-shuffle replicates. Null-extremeness score: derived from the position of observed summary statistics in the null distribution.

Pair-dominant pairs are characterised by very high pairwise Δr values (0.83–0.96) with relatively modest subnetwork-level mean $|\Delta r|$ (~0.15), suggesting that the co-expression change is largely localised to the gene pair itself rather than being diffusely distributed across the subnetwork. Module-dominant pairs also show high pairwise Δr values (0.84–0.86) but substantially elevated subnetwork mean $|\Delta r|$ (~0.32–0.36), consistent with these pairs residing at the core of a more extensively reorganised local co-expression module. This distinction carries direct implications for experimental design: pair-dominant pairs are primarily predicted to

exhibit strong focal pair-edge perturbation phenotypes, whereas module-dominant pairs may produce broader network-level effects upon perturbation.

Label-shuffled null-strength values reached the computational ceiling (300) for all 10 primary pairs at both pair and subnetwork levels, indicating that none of the 1,200 label-permuted iterations produced statistics approaching the observed values. This result argues against the possibility that the tumour-specific signals arise from label-dependent bias in the analytic procedure.

3.6 Null-integrated priority panel and lead candidate identification

Integration of all validation layers into a null-integrated priority table yielded an eight-pair focus panel (three pair-dominant, three module-dominant, two mixed) with overall priority scores ranging from 0.416 to 0.593 (Table 4). The highest-ranked pair (ENSG00000183305.13 || ENSG00000268606.5) exhibited an overall priority score of 0.593, pair-dominant score of 0.815, null-extremeness score of 32.9, and pair-to-module ratio of 6.35, placing it distinctly above the second-ranked pair (overall score 0.461; Δscore = 0.132 between ranks 1 and 2 in the null-integrated table, and 0.269 in the evidence-refreshed table).

Table 4. Null-integrated priority scores and architectural classification for the eight-pair focus panel.

PAIR (ENSEMBL IDS)	SUBTYPE	PAIR SCORE	MODULE SCORE	OVERALL SCORE	NULL EXTREMENESS	PAIR/MODULE RATIO	REFRESHED PRIORITY	DECI
ENSG00000183305.13 || ENSG00000268606.5	Pair-dominant	0.815	0.372	0.593	32.9	6.35	0.653	**Adva**
ENSG00000197172.10 || ENSG00000221867.8	Pair-dominant	0.553	0.369	0.461	36.3	5.66	0.384	Rese
ENSG00000183305.13 || ENSG00000221867.8	Pair-dominant	0.530	0.359	0.444	32.6	5.63	0.350	Rese
ENSG00000201098.1 || ENSG00000206834.1	Module-dominant	0.258	0.671	0.464	8.1	2.58	0.241	Depr
ENSG00000201098.1 || ENSG00000276027.1	Mixed	0.323	0.602	0.463	8.4	2.77	0.238	Depr
ENSG00000200156.1 || ENSG00000207765.1	Module-dominant	0.234	0.634	0.434	9.8	2.56	0.189	Depr
ENSG00000200156.1 || ENSG00000277300.1	Mixed	0.238	0.625	0.432	8.4	2.64	0.186	Depr
ENSG00000201098.1 || ENSG00000207765.1	Module-dominant	0.233	0.599	0.416	9.3	2.58	0.159	Depr

Refreshed priority incorporates first-pass external evidence annotation (LUAD literature relevance, dependency support, synthetic lethality/vulnerability, novelty status, actionability, evidence confidence). "Advance": primary experimental candidate; "Reserve":

Threshold sensitivity analysis confirmed that the lead candidate (ENSG00000183305.13 || ENSG00000268606.5) remains the sole pair with refreshed priority score exceeding 0.40 across all tested thresholds (0.35–0.65), demonstrating that its priority advantage is not contingent on a specific scoring parameterisation. A rescue analysis simulating the effect of upgrading low-confidence external evidence annotations yielded a maximum potential gain of 0.048 score units for the highest-ranked deprioritised pair, which would not be sufficient to threaten the lead candidate's rank position.

3.7 Batch A experimental design framework

For the eight-pair focus panel, a structured perturbation experimental design was generated comprising 32 experimental rows (four arms per pair: gene-A-only, gene-B-only, A+B combined, and unperturbed control). Track-specific primary hypotheses were defined: for pair-dominant pairs, the prediction is that A+B combined perturbation will more effectively disrupt the tumour-stabilised focal co-expression edge than either single perturbation, with preferential tumour selectivity; for module-dominant pairs, the prediction is that A+B combined perturbation will produce greater local subnetwork destabilisation than single perturbations, with tumour-selective phenotypic consequences. The lead candidate was assigned Tier 1 / Phase 1A priority; six further pairs received Tier 2 / Phase 2 status; one pair was assigned Tier 3 / Phase 3. Suggested tumour model count per Tier 1 pair: three tumour models and two normal comparator models.

4. DISCUSSION

We present a large-scale systematic analysis of pairwise co-expression instability in LUAD that covers, to our knowledge, all gene pairs within a validated 4,000-gene joint tumour–normal universe (7,998,000 pairs total). Several findings warrant discussion.

The observation that primary candidate pairs exhibit perfect sign consistency across 200 independent resampling iterations (sign_consistency = 1.00 for all 10 pairs) is not trivially expected in RNA-seq transcriptomics data, which are characterised by substantial sample-to-sample variability, platform-specific noise, and batch effects. The persistence of this stability under 80% subsampling of a 513-sample tumour cohort suggests that the underlying co-expression strengthening is a robust, widespread property of LUAD tumour transcriptomes rather than a feature driven by a minority subgroup or outlier samples.

The local network context analysis provides a dimension of evidence that is conceptually independent from the pairwise correlation magnitude itself. The finding that primary pairs acquire shared tumour-specific co-expression neighbourhoods—detectable through Jaccard overlap of each gene's top-25 partners in the tumour matrix—suggests that these pairs do not merely undergo bilateral correlation strengthening in isolation, but become embedded within tumour-specific local co-regulatory modules. This observation is consistent with the concept of module-level transcriptional reorganisation in cancer proposed by network medicine frameworks.[7,9] The complete absence of such shared neighbourhoods in normal tissue (mean Jaccard = 0.000 for primary pairs in normal) implies that the co-regulatory relationship is tumour-acquired rather than being a constitutively strong but variable association.

The identification of two architecturally distinct primary pair subtypes has implications for both experimental design and biological interpretation. Pair-dominant pairs, characterised by high pair-to-module ratios

(~5.3–6.4) and relatively low subnetwork-level mean $|\Delta r|$ (~0.15), may represent focused relational breakpoints at which tumour transcriptional reorganisation is concentrated within a single co-expression edge. These are mechanistically analogous to the strong focal connectivity changes previously associated with transcription factor–target or co-factor interactions that are specifically altered in tumour contexts.[10] Module-dominant pairs, in which the 800-gene subnetwork mean $|\Delta r|$ is approximately twofold higher (~0.32–0.36), may instead mark the centre of gravity of a more extensively reorganised local co-expression module, analogous to network hubs within tumour-specific transcriptional regulons. The implication for experimental follow-up is that the success criteria, readout metrics and perturbation scales appropriate for these two subtypes may differ substantially, and that their conflation within a single experimental paradigm risks false-negative conclusions for the module-dominant class.

The null-shuffled permutation analysis, which confirms that primary pair signals were never exceeded by 1,200 label-permuted null iterations in either pair-level or subnetwork-level metrics, provides a form of empirical specificity that complements the resampling stability analysis. Whereas resampling addresses the question of whether the signal is stable under sampling variation within the observed labels, label shuffling addresses whether the observed label assignment produces signals that are distinguishable from random. The saturation of null-strength metrics (reaching a computational ceiling of 300) indicates that the null permutation test did not encounter a single replicate approaching the observed statistics for any of the 10 primary pairs, suggesting the signals are highly specific to the actual tumour–normal stratification.

Several limitations of this study require acknowledgement. First, the external evidence annotations used in the final evidence-integrated priority scoring are first-pass, manually assigned placeholders rather than systematic curations from LUAD-specific literature or dependency databases. Consequently, the deprioritisation of module-dominant and mixed pairs in the final decision layer reflects the current absence of external supporting evidence rather than a biological determination of non-relevance; these pairs may be appropriately re-evaluated following targeted literature and database curation. Second, the reported gene identifiers are Ensembl IDs; gene name annotation and pathway assignment are pending and will be presented in a subsequent, experimentally validated report. Third, the analysis is purely correlational and observational at the mRNA expression level; it does not establish causal or mechanistic relationships between identified gene pairs and tumour biology. The functional perturbation design described (Section 2.10) is intended to address this gap in the next experimental phase.

The lead candidate pair (ENSG00000183305.13 || ENSG00000268606.5) presents several properties that collectively motivate its prioritisation for experimental follow-up: the highest overall null-integrated priority score (0.593), the highest pair-dominant score (0.815), the highest pair-to-module ratio (6.35), a refreshed priority score that remains uniquely above 0.40 across all tested thresholds, and a score gap of 0.269 relative to the second-ranked pair—a gap that is not reducible by plausible improvements in external evidence annotation for lower-ranked pairs. These properties are necessary but not sufficient conditions for experimental efficacy; confirmation through functional perturbation in tumour versus normal cellular models is required before any therapeutic utility can be proposed.

5. CONCLUSIONS

We have systematically characterised pairwise co-expression differential signals across 7,998,000 gene pairs in LUAD tumour versus normal tissue and identified a primary candidate panel of 10 gene pairs with consis-

tent, robust, tumour-specific co-expression strengthening. These pairs exhibit tumour-localised local network coherence, resistance to label-permutation-based null testing, and stratify into two architecturally distinct subtypes with distinct predicted perturbation phenotypes. An evidence-integrated priority framework identifies one primary experimental candidate, two reserve candidates for targeted evidence review, and five currently deprioritised pairs. The computational framework described is fully checkpointed, batch-parallelised, and reproducible, and may be applicable to systematic pairwise instability analysis in other tumour types. Experimental validation of the perturbation hypotheses derived from this analysis is the designated next step.

6. DECLARATIONS

Data availability

TCGA-LUAD expression data are publicly available via the GDC portal (https://portal.gdc.cancer.gov). Derived analysis outputs, including the prioritised pair table and robustness metrics, will be available on Zenodo.

REFERENCES

[1] Siegel RL, Miller KD, Wagle NS, Jemal A. Cancer statistics, 2023. *CA Cancer J Clin.* 2023;73(1):17–48.

[2] Planchard D, Popat S, Kerr K, et al. Metastatic non-small cell lung cancer: ESMO Clinical Practice Guidelines for diagnosis, treatment and follow-up. *Ann Oncol.* 2018;29(Suppl 4):iv192–iv237.

[3] Barabási AL, Gulbahce N, Loscalzo J. Network medicine: a network-based approach to human disease. *Nat Rev Genet.* 2011;12(1):56–68.

[4] Vidal M, Cusick ME, Barabási AL. Interactome networks and human disease. *Cell.* 2011;144(6):986–998.

[5] Ideker T, Sharan R. Protein networks in disease. *Genome Res.* 2008;18(4):644–652.

[6] Jerby-Arnon L, Pfetzer N, Waldman YY, et al. Predicting cancer-specific vulnerability via data-driven detection of synthetic lethality. *Cell.* 2014;158(5):1199–1209.

[7] Califano A, Alvarez MJ. The recurrent architecture of tumour initiation, progression and drug sensitivity. *Nat Rev Cancer.* 2017;17(2):116–130.

[8] Carro MS, Lim WK, Alvarez MJ, et al. The transcriptional network for mesenchymal transformation of brain tumours. *Nature.* 2010;463(7279):318–325.